

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 031 928 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
30.08.2000 Bulletin 2000/35

(51) Int Cl.7: **G06F 11/10**

(21) Application number: **00111252.3**

(22) Date of filing: **02.06.1994**

(84) Designated Contracting States:
**AT BE CH DE DK ES FR GB GR IE IT LI LU MC NL
PT SE**
Designated Extension States:
SI

- **Malcolm, Michael**
California 94022 (US)
- **Lau, James**
Cupertino, CA 95014 (US)
- **Rakitzis, Byron**
130 Mountain View, CA 94043 (US)

(30) Priority: **04.06.1993 US 71798**

(62) Document number(s) of the earlier application(s) in
accordance with Art. 76 EPC:
94919367.6 / 0 701 715

(71) Applicant: **Network Appliance, Inc.**
Sunnyvale, California 94089 (US)

(74) Representative: **Leeming, John Gerard**
J.A. Kemp & Co.,
14 South Square,
Gray's Inn
London WC1R 5LX (GB)

(72) Inventors:
• **Hitz, David**
Portola Valley, CA 94028 (US)

Remarks:

This application was filed on 25 - 05 - 2000 as a
divisional application to the application mentioned
under INID code 62.

(54) **A method for providing parity in a raid sub-system using non-volatile memory**

(57) The present invention is a method for providing
error correction for an array of disks (828) using non-
volatile random access memory (NV-RAM) (816). Non-
volatile RAM (816) is used to increase the speed of
RAID recovery from a disk error(s). This is done by list-
ing of all blocks for which the parity is possibly incons-
istent. This list of blocks is much smaller than the total
number of parity blocks (820) in the RAID. The total

number of parity blocks (820) in the RAID (828) is in the
range of hundreds of thousands. Knowledge of the
number of parity blocks that are possibly inconsistent
makes it possible to fix only those few blocks, identified
in the list, resulting in a significant time savings. The
technique for safely writing to a RAID with a broken disk
is complicated. In this technique, data that can become
corrupted is copied into NV-RAM (816) before the po-
tentially corrupting operation is performed.

EP 1 031 928 A2

Description

[0001] The present invention is related to the field of error correction techniques for an array of disks.

[0002] A computer system typically requires large amounts of secondary memory, such as a disk drive, to store information (e.g. data and/or application programs). Prior art computer systems often use a single "Winchester" style hard disk drive to provide permanent storage of large amounts of data. As the performance of computers and associated processors has increased, the need for disk drives of larger capacity, and capable of high speed data transfer rates, has increased. To keep pace, changes and improvements in disk drive performance have been made. For example, data and track density increases, media improvements, and a greater number of heads and disks in a single disk drive have resulted in higher data transfer rates.

[0003] A disadvantage of using a single disk drive to provide secondary storage is the expense of replacing the drive when greater capacity or performance is required. Another disadvantage is the lack of redundancy or back up to a single disk drive. When a single disk drive is damaged, inoperable, or replaced, the system is shut down.

[0004] One prior art attempt to reduce or eliminate the above disadvantages of single disk drive systems is to use a plurality of drives coupled together in parallel. Data is broken into chunks that may be accessed simultaneously from multiple drives in parallel, or sequentially from a single drive of the plurality of drives. One such system of combining disk drives in parallel is known as "redundant array of inexpensive disks" (RAID). A RAID system provides the same storage capacity as a larger single disk drive system, but at a lower cost. Similarly, high data transfer rates can be achieved due to the parallelism of the array.

[0005] RAID systems allow incremental increases in storage capacity through the addition of additional disk drives to the array. When a disk crashes in the RAID system, it may be replaced without shutting down the entire system. Data on a crashed disk may be recovered using error correction techniques.

RAID Arrays

[0006] RAID has six disk array configurations referred to as RAID level 0 through RAID level 5. Each RAID level has advantages and disadvantages. In the present discussion, only RAID levels 4 and 5 are described. However, a detailed description of the different RAID levels is disclosed by Patterson, et al. in *A Case for Redundant Arrays of Inexpensive Disks (RAID)*, ACM SIGMOD Conference, June 1988. This article is incorporated by reference herein.

[0007] RAID systems provide techniques for protecting against disk failure. Although RAID encompasses a number of different formats (as indicated above), a common feature is that a disk (or several disks) stores parity information for data stored in the array of disks. A RAID level 4 system stores all the parity information on a single parity disk, whereas a RAID level 5 system stores parity blocks throughout the RAID array according to a known pattern. In the case of a disk failure, the parity information stored in the RAID subsystem allows the lost data from a failed disk to be recalculated.

[0008] Figure 1 is a block diagram illustrating a prior art system implementing RAID level 4. The system comprises N+1 disks 112-118 coupled to a computer system, or host computer, by communication channel 130. In the example, data is stored on each hard disk in 4 KByte (KB) blocks or segments. Disk 112 is the Parity disk for the system, while disks 114-118 are Data disks 0 through N-1. RAID level 4 uses disk "striping" that distributes blocks of data across all the disks in an array as shown in Figure 1. A stripe is a group of data blocks where each block is stored on a separate disk of the N disks along with an associated parity block on a single parity disk. In Figure 1, first and second stripes 140 and 142 are indicated by dotted lines. The first stripe 140 comprises Parity 0 block and data blocks 0 to N-1. In the example shown, a first data block 0 is stored on disk 114 of the N+1 disk array. The second data block 1 is stored on disk 116, and so on. Finally, data block N-1 is stored on disk 118. Parity is computed for stripe 140 using well-known techniques and is stored as Parity block 0 on disk 112. Similarly, stripe 142 comprising N data blocks is stored as data block N on disk 114, data block N+1 on disk 116, and data block 2N-1 on disk 118. Parity is computed for the 4 stripe 142 and stored as parity block 1 on disk 112.

[0009] As shown in Figure 1, RAID level 4 adds an extra parity disk drive containing error-correcting information for each stripe in the system. If an error occurs in the system, the RAID array must use all of the drives in the array to correct the error in the system. RAID level 4 performs adequately when reading small pieces of data. However, a RAID level 4 array always uses the dedicated parity drive when it writes data into the array.

[0010] RAID level 5 array systems also record parity information. However, it does not keep all of the parity sectors on a single drive. RAID level 5 rotates the position of the parity blocks through the available disks in the disk array of N+1 disk. Thus, RAID level 5 systems improve on RAID 4 performance by spreading parity data across the N+1 disk drives in rotation, one block at a time. For the first set of blocks, the parity block might be stored on the first drive. For the second set of blocks, it would be stored on the second disk drive. This is repeated so that each set has a parity block, but not all of the parity information is stored on a single disk drive. In RAID level 5 systems, because no single disk holds all of the parity information for a group of blocks, it is often possible to write to several different drives in the

array at one instant. Thus, both reads and writes are performed more quickly on RAID level 5 systems than RAID 4 array.

[0011] Figure 2 is a block diagram illustrating a prior art system implementing RAID level 5. The system comprises N+1 disks 212-218 coupled to a computer system or host computer 120 by communication channel 130. In stripe 240, parity block 0 is stored on the first disk 212. Data block 0 is stored on the second disk 214, data block 1 is stored on the third disk 216, and so on. Finally, data block N-1 is stored on disk 218. In stripe 212, data block N is stored on the first disk 212. The second parity block 1 is stored on the second disk 214. Data block N+1 is stored on disk 216, and so on. Finally, data block 2N-1 is stored on disk 218. In M-1 stripe 244, data block MN-N is stored on the first disk 212. Data block MN-N+1 is stored on the second disk 214. Data block MN-N+2 is stored on the third disk 216, and so on. Finally, parity block M-1 is stored on the nth disk 218. Thus, Figure 2 illustrates that RAID level 5 systems store the same parity information as RAID level 4 systems, however, RAID level 5 systems rotate the positions of the parity blocks through the available disks 212-218.

[0012] In RAID level 5, parity is distributed across the array of disks. This leads to multiple seeks across the disk. It also inhibits simple increases to the size of the RAID array since a fixed number of disks must be added to the system due to parity requirements.

[0013] The prior art systems for implementing RAID levels 4 and 5 have several disadvantages. The first disadvantage is that, after a system failure, the parity information for each stripe is inconsistent with the data blocks stored on the other disks in the stripe. This requires the parity for the entire RAID array to be recalculated. The parity is recomputed entirely because there is no method for knowing which parity blocks are incorrect. Thus, all the parity blocks in the RAID array must be recalculated. Recalculating parity for the entire RAID array is highly time consuming since all of the data stored in the RAID array must be read. For example, reading an entire 2 GB disk at maximum speed takes 15 to 20 minutes to complete. However, since few computer systems are able to read very many disks in parallel at maximum speed, recalculating parity for a RAID array takes even longer.

[0014] One technique for hiding the time required to recompute parity for the RAID array is to allow access to the RAID array immediately, and recalculate parity for the system while it is on-line. However, this technique suffers two problems. The first problem is that, while recomputing parity, blocks having inconsistent parity are not protected from further corruption. During this time, a disk failure in the RAID array results in permanently lost data in the system. The second problem with this prior art technique is that RAID subsystems perform poorly while calculating parity. This occurs due to the time delays created by a plurality of input/output (I/O) operations imposed to recompute parity.

[0015] The second disadvantage of the prior art systems involves writes to the RAID array during a period when a disk is not functioning. Because a RAID subsystem can recalculate data on a malfunctioning disk using parity information, the RAID subsystem allows data to continue being read even though the disk is malfunctioning. Further, many RAID systems allow writes to continue although a disk is malfunctioning. This is disadvantageous since writing to a broken RAID array can corrupt data in the case of a system failure. For example, a system failure occurs when an operating system using the RAID array crashes or when a power for the system fails or is interrupted otherwise. Prior art RAID subsystems do not provide protection for this sequence of events.

[0016] The present invention is defined in the appended claims and provides error correction for an array of disks using non-volatile random access memory (NV-RAM).

[0017] Non-volatile RAM is used to increase the speed of RAID recovery from disk error(s). This is accomplished by keeping a list of all disk blocks for which the parity is possibly inconsistent. Such a list of disk blocks is smaller than the total number of parity blocks in the RAID subsystem. The total number of parity blocks in the RAID subsystem is typically in the range of hundreds of thousands of parity blocks. Knowledge of the number of parity blocks that are possibly inconsistent makes it possible to fix only those few blocks, identified in the list, in a significantly smaller amount of time than is possible in the prior art. The present invention also provides a technique of protecting against simultaneous system failure and a broken disk and of safely writing to a RAID subsystem with one broken disk.

[0018] The present invention will be described further below with reference to the following description of exemplary embodiments and the accompanying drawings, in which:

Figure 1 is a block diagram of a prior art RAID level 4 subsystem;

Figure 2 is a block diagram of a prior art RAID level 5 subsystem;

Figures 3A-3C are prior art diagrams illustrating recomputation of data stored in a "stripe";

Figures 4A-4B are prior art timing diagrams for parity corruption on system failure;

Figure 5 is a timing diagram of the present invention for preventing data corruption on occurrence of a write to a malfunctioning disk;

Figure 6 is a prior art timing diagram illustrating data corruption on a write with a broken disk;

Figures 7A-7B are timing diagrams of the present invention for preventing data corruption on a write with a broken disk;

Figure 8 is a diagram illustrating the present invention.

Figure 9 is a prior art timing diagram illustrating data corruption for simultaneous system and disk failures;

Figures 10A-10C are timing diagrams of the present invention preventing data corruption for simultaneous system and disk failures; and,

Figures 11A-11C are flow diagrams of the present invention illustrating the process of recovery.

Figures 12A-12C are timing diagrams for parity by recalculation.

[0019] A method and apparatus for providing error correction for an array of disks using non-volatile random access memory (NV-RAM) is described. In the following description, numerous specific details, such as number and nature of disks, disk block sizes, etc., are described in detail in order to provide a more thorough description of the present invention. It will be apparent, however, to one skilled in the art, that the present invention may be practiced without these specific details. In other instances, well-known features have not been described in detail so as not to unnecessarily obscure the present invention.

[0020] In particular, many examples consider the case where only one block in a stripe is being updated, but the techniques described apply equally well to multi-block updates.

[0021] The present invention provides a technique for: reducing the time required for recalculating parity after a system failure; and, preventing corruption of data in a RAID array when data is written to a malfunctioning disk and the system crashes. The present invention uses non-volatile RAM to reduce these problems. A description of the prior art and its corresponding disadvantages follows. The disadvantages of the prior art are described for: parity corruption on a system failure; data corruption on write with broken disk; and, data corruption with simultaneous system and disk failures.

Recomputing Lost Data With RAID

[0022] Parity is computed by Exclusive-ORing the data blocks stored in a stripe. The parity value computed from the N data blocks is recorded in the parity block of the stripe. When data from any single block is lost (i.e., due to a disk failure), the lost data for the disk is recalculated by Exclusive-ORing the remaining blocks in the stripe. In general, whenever a data block in a stripe is modified, parity must be recomputed for the stripe. When updating a stripe by writing all N data blocks, parity can be computed without reading any data from disk and parity and data can be written together, in just one I/O cycle. Thus, writing to all N data blocks in a stripe requires a minimum amount of time. When writing a single data block to disk, parity-by-subtraction is used (described below). One I/O cycle is required to read the old data and parity, and a second I/O cycle is required to write the new data and parity. Because the spindles of the disks in the RAID array are not synchronized, the writes do not generally occur at exactly the same time. In some cases, the parity block will reach the disk first, and in other cases, one of the data blocks will reach the disk first. The techniques described here do not depend on the order in which blocks reach the disk.

[0023] Another alternative for disks having non-synchronized spindles is for parity to be computed first and the parity block written to disk before a data block(s) is written to disk. Each data block on a disk in the RAID array stores 4 KB of data. In the following discussion, the data in each 4 KB block is viewed as a single, large integer (64 K-bits long). Thus, the drawings depict integer values for information stored in the parity and data disk blocks. This convention is used for illustration only in order to simplify the drawings.

[0024] Figure 3A is a diagram illustrating a prior art RAID level 4 subsystem, where $N = 3$, comprising four disks, 330-336. In the diagram, disk 330 is the parity disk. Disks 332-336 are data disks. The diagram illustrates a stripe 320 contained on the disks 330-336 in the RAID array. Disk block 330A is a parity block containing the integer value 12. Disk blocks 332A-336A are data blocks of the stripe 320, respectively. Data blocks 332A-336A contain data values of 4, 7 and 1, respectively. Data for each block 332A-336A in a single stripe 320 is represented as an integer. Parity for stripe 320 is represented as the sum of data values stored in data blocks 332A-336A. Parity block 330A contains the value 12 (i.e., $4+7+1$). Figure 3A is a drawing that is merely one example of an error correction technique using parity. The parity value is the Exclusive-OR of the data blocks 332A-336A, but the mathematical properties of addition match those of the Exclusive-OR function. Therefore, addition is used in Figure 3A.

[0025] Figure 3B is a timing diagram of activity on the stripe 320 illustrated in Figure 3A. The table has headings of Parity, Data 0, Data 1 and Data 2. The values 12, 4, 7 and 1 are illustrated under the corresponding table headings.

[0026] Figure 3B is a table illustrating a stripe having a lost data block at time T_B . As illustrated in Figure 3B, stripe 320 contains lost data in data block 1 from data disk 334 of Figure 3A. This is illustrated in the table by a question mark enclosed in a box under the data 1 heading. At time T_A , parity, data 0 and data 2 have values of 12, 4 and 1, respectively. The data on disk 334 for data block 1 can be recalculated in real time as follows:

$$\text{Data 1} = \text{Parity} - \text{Data 0} - \text{Data 2} = 12 - 4 - 1 = 7, \quad (1)$$

where data block 1 is computed using the parity block, data block 0 and data block 2. Thus, the data value 7 stored in data block 1 of disk 334 shown in Figure 3A can be recomputed at time T_C . In Figure 3B, at time T_C , the value 7 that has been recomputed for data block 1, is indicated by being enclosed within parentheses. In subsequent figures, recomputed values are represented using parentheses. That is, the parentheses indicate data for a broken disk as computed by the parity and data on the other disks.

[0027] As shown in Figure 3B, data on a broken disk can be recomputed using the parity disk and the remaining disks in the disk array. The broken disk 334 of Figure 3A can eventually be replaced and the old contents of the disk can be recalculated and written to a new disk. Figure 3C is a block diagram of the RAID subsystem containing a new data 1 disk 338. As shown in Figure 3E, stripe 320 has values of 12, 4, 7 and 1 for parity, data 0, new data 1 and data 2. These values are stored in parity block 330A and data blocks 332A, 338A and 336A. Thus, a new disk 338 replaces broken disk 334 of the RAID system, and the data value stored previously in data block 334A of disk 334 can be computed as shown above and stored in data block 338A of replacement disk 338.

[0028] When new data is written to a data block, the parity block is also updated. Parity is easily computed, as described above, when all data blocks in a stripe are being updated at once. When this occurs, the new value for parity is recalculated from the information being written to the disks. The new parity and data blocks are then written to disk. When only some of the data blocks in a stripe are modified, updating the parity block is more difficult since more I/O operations are required. There are two methods for updating parity in this case: parity update by subtraction; and, parity update by recalculation.

[0029] For example, when a single data block is written, the RAID system can update parity by subtraction. The RAID system reads the parity block and the block to be overwritten. It first subtracts the old data value from the parity value, adds the new data value of the data block to the intermediate parity value, and then writes both the new parity and data blocks to disk.

[0030] For recalculation of parity, the RAID system first reads the other N-1 data blocks in the stripe. After reading the N-1 data blocks, the RAID system recalculates parity from scratch using the modified data block and the N-1 data blocks from disk. Once parity is recalculated, the new parity and data blocks are written to disk.

[0031] Both the subtraction and recalculation technique for updating parity can be generalized to situations where more than one data block is being written to the same stripe. For subtraction, the parity blocks and the current contents of all data blocks that are about to be overwritten are first read from disk. For recalculation, the current contents of all data blocks that are not about to be overwritten are first read from disk. The instance where all N data blocks in the stripe are written simultaneously is a degenerate case of parity by recalculation. All data blocks that are not being written are first read from disk, but in this instance, there are no such blocks.

How Stripes Become Inconsistent During System Failure

[0032] An inconsistent stripe comprises a parity block that does not contain the Exclusive-OR of all other blocks in the stripe. A stripe becomes inconsistent when a system failure occurs while some of the writes for an update have been completed but others have not. For example, when a first data block is being overwritten. As previously described, the parity block for the stripe is recomputed and overwritten as well as the data block. When the system fails after one of the data blocks has been written to disk, but not the other, then the stripe becomes inconsistent.

[0033] A stripe can only become inconsistent when it is being updated. Thus, the number of potentially inconsistent stripes at any instant is limited to the number of stripes that are being updated. For this reason, the present invention maintains a list in NV-RAM comprising all the stripes that are currently being updated. Since only these stripes can potentially be corrupted, parity is recalculated after a system failure for only the stripes stored in the list in NV-RAM. This greatly reduces the total amount of time required for recalculating parity after a system failure in comparison to the prior art methods, described previously, that take much longer.

Parity Corruption On A System Failure In The Prior Art

[0034] In the following diagrams, the value indicated within parentheses for a malfunctioning data disk is not an actual value stored on disk. Instead, it is a calculated value retained in memory for the broken disk in the RAID array.

[0035] Figure 4A is a prior art diagram illustrating a system crash while changing values are written to disks 330-336 of Figure 3A. The diagram is for the case where the data block reaches the disk before the parity block. As indicated in Figure 4A, time is increasing in a downward direction. At time T_A , the parity block has a value of 12 and data blocks 0 to 2 have values of 4, 7, and 1, respectively. At time T_B , a new value of 2 is written (indicated by a box around the value 2) to data block 0, thereby replacing the value of 4 that is stored in data block 0 at time T_A . The other values stored in data blocks 1 and 2 do not change. When operating normally, the prior art writes a new parity value of 10 (indicated by a box) at time T_C to the parity disk as indicated under the parity heading. This updates the parity block for the write to data block 0 at time T_B . The new value of 10 for parity at time T_C is computed from the values of 2, 7, and 1 of data blocks 0 to 2, respectively. Thus, the timing diagram in Figure 4A illustrates a prior art RAID subsystem in the case where the data block reaches disk before the parity block.

[0036] When a system failure occurs between time T_B and T_C in Figure 4A, parity is corrupted for the stripe. The timing diagram shows that a new data value of 2 is written to data disk 0 at time T_B before recomputed parity for the stripe is updated. Thus, when the RAID subsystem subsequently restarts, the parity disk has the old value of 12 (indicated by an underline) instead of the correct value of 10. This occurs since the stripe was not updated before the system failure occurred. The parity for the stripe is now corrupted since:

$$\text{Parity} = \text{Data 0} + \text{Data 1} + \text{Data 2} = 2 + 7 + 1 = 10 \neq 12. \quad (2)$$

[0037] Similarly, Figure 4B is another prior art diagram illustrating a system crash while changing values are written to disks 330-336 of Figure 3A. The diagram is for the case where the parity block reaches disk before the data block. At time T_A , the parity block has a value of 12 and data blocks 0 to 2 have values of 4, 7, and 1, respectively. At time T_B , a new value of 10 is written (indicated by a box around the value 2) to the parity block, thereby replacing the value of 12 that is stored in the parity block at time T_A . The data values stored in data blocks 0-2 do not change. The new value of 10 for parity at time T_B is computed from the values of 7 and 1 for data blocks 1 and 2, respectively, and the new value of 2 for data block 0. When operating normally, the prior art writes the new data value of 2 (indicated by a box) at time T_C to the data disk 0, as indicated under the Data 0 heading. This updates the data block 0 in accordance with the write to the parity block at time T_B . Thus, the timing diagram in Figure 4A illustrates a prior art RAID subsystem in the case where the parity block reaches disk before the data block.

[0038] When a system failure occurs between time T_B and T_C in Figure 4B, parity is corrupted for the stripe. The timing diagram shows that the new parity value of 10 is written to the parity disk at time T_B before data block 0 of the stripe is updated. Thus, when the RAID subsystem subsequently restarts, data disk 0 has the old value of 4 (indicated by an underline) instead of the correct value of 2. This occurs because the stripe was not updated before the system failure occurred. The parity for the stripe is now corrupted since:

$$\text{Parity} = \text{Data 0} + \text{Data 1} + \text{Data 2} = 4 + 7 + 1 = 12 \neq 10. \quad (3)$$

[0039] Figures 4A-4B illustrate two cases of writing new data to a data block and updating the parity disk where the spindles of the disks in the RAID array are not synchronized. The first case shown in Figure 4A illustrates a new data value reaching the data disk first, and then subsequently updating the parity value on the parity disk. The second case illustrated in Figure 4B illustrates parity reaching disk first followed by the data update. For Figures 4A and 4B, when the system fails between times T_B and T_C , corruption of the file system occurs. If the system fails after time T_B in Figures 4A and 4B, then the parity values illustrated are not correct for the system. In the case of the system illustrated in Figure 4A, the new data values have a sum of 10, which is equal to the values of 2, 7 and 1. However, the parity value at time T_B indicates a value of 12. Thus, the parity value stored on the parity disk does not equal the new parity value for the data values stored on data disk 0-2. Similarly, if a failure occurs after time T_B for the second system illustrated in Figure 4B, the data disks 0-1 have values of 4, 7 and 1, respectively. The parity value for these data blocks is equal to 12. However, parity in this system is first updated before writing the new data value to disk, therefore, the parity stored on the parity disk at time T_B is equal to 10. Thus, subsequent to time T_B , the parity stored on the parity disk does not equal the parity value for the data blocks since the new data was not updated before the system failed.

[0040] In the prior art, after a system fails, parity is recalculated for all of the stripes on occurrence of a system restart. This method of recalculating parity after a failure for all stripes requires intensive calculations, and therefore, is very slow. The present invention is a method for recalculating parity after a system failure. The system maintains a list of

stripes having writes in progress in non-volatile RAM. Upon restarting after a system failure, just the list of stripes with writes in progress that are stored in non-volatile RAM are recalculated.

Data Corruption On Write With Broken Disk In The Prior Art

5

[0041] When writing to a RAID array that has a malfunctioning or broken disk, data corruption occurs during system failure. Figure 6 is a prior art diagram illustrating data corruption for a malfunctioning disk when a system failure occurs where the data disk is updated for the new data value before parity is written to disk. In Figure 6, data disk 1 is shown to be malfunctioning by indicating values within parentheses under the Data 1 heading. At time T_A , the parity disk has a value of 12. Prior to time T_A when data disk 1 malfunctions, the parity disk value is equal to the sum of data disks 0 to 2 having values of 4, 7 and 1, respectively. The value of 7 for data block 1 at time T_A is enclosed within parentheses. This value does not represent a value stored on data disk 1, but instead is computed from the parity block and data blocks 0 and 2 of the stripe as follows:

15

$$\text{Data 0} = \text{Parity} - \text{Data 1} - \text{Data 2} = 12 - 4 - 1 = 7. \quad (4)$$

[0042] At time T_B , a new value of 2 is written to data disk 0 (indicated by enclosing 2 within a box). At time T_B , parity has not been updated for the new value of 2 written to data disk 0 and has a value of 12. Thus, the computed value for data block 1 is 9 instead of 7. This is indicated in Figure 6 by enclosing the value 9 within parentheses for data disk 1 at time T_B .

20

[0043] When operating normally at time T_C , the parity block is updated to 10 due to the value of 2 written to data block 0 at time T_B . The new value of 10 for parity at time T_C is indicated within a rectangle. For a parity value of 10, the correct value of 7 for data block 1 is indicated within parentheses. As indicated in the Figure 6, because data disk 1 is broken, the data stored in data block 1 is calculated based on the other blocks in the disk array. After the first write at time T_B for data block 0, the computed value of data block 1 is incorrect. The value of 9 for data block 1 is incorrect until the second write for parity at time T_C is completed.

25

[0044] When a system failure occurs between times T_B and T_C , writing to a RAID array that has a malfunctioning or broken disk corrupts data in the stripe. As shown in Figure 6 for the prior art, parity is not updated and therefore has a value of 12 (indicated by an underline). Thus, the computed value for data block 1 of the stripe is incorrect and the stripe is corrupted as follows:

30

$$\text{Data 1} = \text{Parity} - \text{Data 0} - \text{Data 2} = 12 - 2 - 1 = 9 \neq 7. \quad (5)$$

35

[0045] Similar corruption of data occurs for the case where parity reaches disk before data does.

Data Corruption With Simultaneous System And Disk Failures

[0046] RAID systems are most likely to experience a disk failure when a system failure occurs due to power interruption. Commonly, a large, transient voltage spike occurring after power interruption damages a disk. Thus, it is possible for a stripe to be corrupted by simultaneous system and disk failures.

40

[0047] Figure 9 is a prior art diagram illustrating simultaneous system and disk failures where the data disk is updated for a new data value before parity is written to disk. At time T_A , the parity disk has a value of 12 and data disks 0-2 have values of 4, 7, and 1, respectively. At time T_B , a new value of 2 is written to data disk 0 (indicated by a box). At time T_B , parity is not updated for the new value of 2 written to data disk 0 and has a value of 12. When a system failure occurs between times T_B and T_C , the value of disk 1 is corrupted. This occurs due to simultaneous system and disk failures between times T_B and T_C .

45

[0048] At time T_C , parity is not updated due to the system failure and therefore has a value of 12 instead of 10. Further, data disk 1 is corrupted due to the disk failure. The computed value of 9 for data block 1 is incorrect. It is computed incorrectly for data disk 1 using the corrupt parity value as follows:

50

$$\text{Data 1} = \text{Parity} - \text{Data 0} - \text{Data 2} = 12 - 2 - 1 = 9 \neq 7. \quad (7)$$

55

[0049] Data is similarly corrupted for the case where parity reaches disk before data.

Overview Of The Present Invention

[0050] Figure 8 is a diagram illustrating the present invention for providing error correction using NV-RAM for a RAID system comprising host computer 810, RAID controller 814 including NV-RAM 816, and N+1 disks 820-826. Host computer 810 is coupled to RAID controller 814 by a first communications channel 812. RAID controller 814 comprises NV-RAM 816 for storing stripes of the RAID array 828 that are possibly in an inconsistent state. RAID controller 814 is coupled to the N+1 disks 820-826 of the RAID array 828 by a second communications channel 818. The RAID array 828 comprises parity disk 820 and N data disks 822-826, respectively.

[0051] NV-RAM 816 is used to increase the speed of RAID recovery after a system failure by maintaining a list of all parity blocks stored on parity disk 820 that are potentially inconsistent. Typically, this list of blocks is small. It may be several orders of magnitude smaller than the total number of parity blocks in the RAID array 828. For example, a RAID array 828 may comprise hundreds of thousands of parity blocks while the potentially inconsistent blocks may number only several hundred or less. Knowledge of the few parity blocks that are potentially inconsistent facilitates rapid recalculation of parity, since only those parity blocks have to be restored.

[0052] The present invention also uses NV-RAM 816 to safely write data to a RAID array 828 having a broken disk without corrupting data due to a system failure. Data that can be corrupted is copied into NV-RAM 816 before a potentially corrupting operation is performed. After a system failure, the data stored in NV-RAM 816 is used to recover the RAID array 828 into a consistent state.

[0053] Figures 11A-C are flow diagrams illustrating the steps performed by the present invention. Referring first to Figure 11A, a boot operation is executed. At decision block 1101 a check is made to determine if the system has just failed. If decision block returns true, the system proceeds to step 1102 (Figure 11B) and executes a recovery process. If decision block returns false, the system proceeds to step 1103 (Figure 11C) for normal operation.

[0054] Referring now to Figure 11B, a flow diagram of recovery after system failure is illustrated. At decision block 1104 a check is made to determine if there are any remaining stripes specified in the list of dirty stripes in NV-RAM. If decision block 1104 returns false, the system proceeds to a return step. If decision block 1104 returns true, the system proceeds to step 1105.

[0055] At step 1105, the stripe number is obtained. At step 1106, the data blocks of the identified stripe required to recompute parity are read. Parity is recomputed for the stripe at step 1107. At step 1108, the new parity block for the stripe is written. The system then returns to decision block 1104.

[0056] Normal operation is illustrated in Figure 11C. At step 1109, all blocks required to update a stripe are read. At step 1110 a new parity for the stripe is computed using the new data. The stripe number of that stripe is written to the list of dirty stripes in NV-RAM at step 1111. At step 1112, the disk blocks required to update the stripe are written to disk. At step 1113 the number of the stripe is removed from the list of dirty stripes in NV-RAM.

Parity Corruption For A System Failure Using NV-RAM

[0057] Figure 5 is a diagram illustrating the present invention for preventing corruption of parity using NV-RAM. The timing diagram is described with reference to Figures 11A-C. Figure 5 is a timing diagram for the present invention illustrating a system crash while changing values are written to disks 330-336 of Figure 3A. The diagram is for the case where parity is computed for a new data value and the data is written to disk prior to updating the parity block. At time T_A , the parity block has a value of 12 and data blocks 0 to 2 have values of 4, 7, and 1, respectively. At time T_A , step 1109 of Figure 11C for the present invention is performed where blocks necessary to update the stripe are read into memory. The system performs step 1110 where a new parity value is computed dependent upon the new data value (s). At time T_B , step 1111 of Figure 11C is performed where the stripe number is written to NV-RAM.

[0058] At time T_C , in step 1112, the new data value of 2 is written (indicated by a box around the value 2) to data block 0, thereby replacing the value of 4 that is stored in data block 0 at time T_B . The other values stored in data blocks 1 and 2 do not change. First, consider the normal case where the system does not fail. The present invention writes a new parity value of 10 (indicated by a box under the parity heading) at time T_D in step 1112. This updates the parity block for the write to data block 0 at time T_C . At time T_E , in step 1113, the stripe number in NV-RAM is cleared. Thus, the stripe comprising the blocks for the parity disk and data disks 0-2 have values of 10, 2, 7, and 1, respectively.

[0059] Next, consider the case when the system does fail between time t_f and t_D (between steps 1111 and 1113). The system reboots, and begins execution at START in Figure 11A.

[0060] In decision block 1101, at time T_D , when a system fault occurs, decision block 1101 returns true (Yes). The stripe has a value of 12 (indicated by an underline) for parity and values for data disks 0-2 of 2, 7, and 1, respectively. As illustrated in Figure 5 for time T_C , parity is corrupted after a system failure since:

$$\text{Parity} = \text{Data 0} + \text{Data 1} + \text{Data 2} = 2 + 7 + 1 = 10 \neq 12.$$

(9)

However, the stripe can be recovered to a consistent state. NV-RAM includes an indication of the stripes that are candidates for recovery, i.e. a list of stripes that are being updated. Everything but the parity value is available on disk (the "2" having been written to disk at time T_C). The data values for the stripe are read from disk and a new parity value of 10 is calculated.

$$\text{Parity} = \text{Data 0} + \text{Data 1} + \text{Data 2} = 2 + 7 + 1 = 10. \quad (10)$$

Thus, the newly calculated parity value of 10 is written to the parity disk in step 1108 at time T_D , and the stripe is no longer corrupt.

[0061] The following is an example of pseudo code that describes the operation of Figure 11C:

- (1) Read all disk blocks required to update stripe.
- (2) Calculate new parity contents.
- (3) Add stripe # for stripe being written to NV-RAM dirty stripe list.
- (4) Write all disk blocks required to update stripe.
- (5) Remove stripe # for stripe just written from NV-RAM dirty stripe list.

[0062] After a system failure, a part of the start-up procedure of Figure 11B can be described by the following pseudo code:

```

for (all stripes specified in the NV-RAM dirty stripe list)
{
    (1) Read all data blocks in the stripe.
    (2) Recompute the parity block for the stripe.
    (3) Write the new parity block for the stripe.
}

```

Thus, the present invention prevents parity corruption after a system failure by using NV-RAM.

Parity Corruption Detection with a Bitmap Instead of a List

[0063] The previous section describes a technique in which a list of potentially corrupted stripes is kept in NV-RAM so that on reboot after a system failure, only the stripes in the list need to have their parity blocks recalculated. An alternate embodiment of the present invention uses a bitmap in NV-RAM to indicate the potentially corrupted stripes whose parity blocks must be recalculated after a system failure.

[0064] This technique uses a bitmap in which each bit represents a group of one or more stripes. A typical disk array might have 250,000 stripes. If each entry in the bitmap represents a single stripe, the bitmap will be about 32 KB. Letting each bit represent a group of 32 adjacent stripes reduces the size to 1 KB.

[0065] After a system failure, this technique is essentially identical to the "list of stripes" technique, except that the bitmap is used to determine which stripes need parity recalculation instead of the list. All stripes in groups whose bit is set in the bitmap have their parity recalculated.

[0066] Managing the bitmap during normal operation is slightly different than managing the list. It is no longer possible to clear a stripe's entry as soon as the update is complete, because a single bit can indicate activity in more than one stripe. One stripe's update may be done, but another stripe sharing the same bit may still be active.

[0067] Instead, the appropriate bit for a stripe is set just before the stripe is updated, but it is not cleared after the update is complete. Periodically, when the bitmap has accumulated too many entries, all blocks are flushed to disk, ensuring that there can be no inconsistent stripes, and the entire bitmap is cleared. The following pseudo-code implements this:

- (1) Read all blocks required to update stripe.
- (2) Calculate new parity contents.
- (3) Set bitmap entry for stripe being updated.
- (4) Write all disk blocks required to update stripe.
- (5) If bitmap is too full, wait for all blocks to reach disk and clear the entire bitmap.

[0068] In case of system failure, the bitmap results in more blocks to clean than the list, but the savings are still considerable compared with recomputing parity for all stripes in the system. A typical RAID system has 250,000 stripes, so even if 2,500 potentially-corrupted stripes are referenced in the bitmap, that is just 1% of the stripes in the system.

[0069] The bitmap technique is especially useful with write-caching disks which don't guarantee that data will reach disk in the case of power failure. Such disks may hold data in RAM for some period before actually writing it. This means that parity corruption is still a possibility even after the stripe update phase has completed. The list technique would not work, because the stripe's parity is still potentially corrupted even though the stripe has been removed from the list.

[0070] Thus, using the bitmap technique and instructing each disk to flush its internal cache at the same time that the bitmap is cleared, allows the invention to work in combination with write-caching disk drives.

Data Corruption On Write With Broken Disk Using NV-RAM

[0071] The present invention solves this problem for data corruption on occurrence of a write with a malfunctioning disk by saving data from the broken disk in non-volatile RAM. Figure 7A is a timing diagram of the present invention for preventing data corruption by storing data from a malfunctioning disk in NV-RAM. The drawing is discussed with reference to Figures 11A-C. In Figure 7A, data is written to disk before parity is updated. At time T_A , broken data disk 1 is illustrated having a value of 7 indicated within parentheses. The value of 7 within parentheses indicates that data disk 1 is malfunctioning and that it is the computed value for the disk. This value is computed by subtracting the values of 4 and 1 of data disks 0 and 2 from the value of 12 stored in the parity disk. In step 1109, the stripe is read from the RAID array at time T_A . The NV-RAM is erased. This is indicated in Figure 7A by a question mark under the heading for NV-RAM.

[0072] At time T_B , a value of 7 for the malfunctioning data disk 1 is written into NV-RAM according to step 1109. The value of 7 for data disk 1 that is written into NV-RAM is indicated by a rectangular box in Figure 7A. The system then computes a new value of parity for the stripe in step 1110 of Figure 11C.

[0073] At time T_C , a new value of 2 (indicated by a box) for data disk 0 is written to the disk before parity for the stripe is updated according to step 1112. Therefore, at time T_C , the value for data disk 1 is 9 and is indicated within parentheses accordingly. In the normal case, where the system does not fail, a new parity value of 10 is written to disk at time T_D , and the computed value of disk 1 becomes 7 again, which is correct. When a system failure occurs between times T_C and T_D , a new value of parity is updated correctly using NV-RAM with respect to the value of 2 written to data disk 0 at time T_C .

[0074] The parity is correctly updated at time T_D by first reading the value for all functioning data disks, according to step 1106, stored in NV-RAM, and recalculating its value as follows:

$$\text{Parity} = \text{Data 0} + \text{NV - RAM} + \text{Data 2} = 2 + 7 + 1 = 10. \quad (12)$$

Thus, a correct value of 10 is computed for parity when the present invention restarts after a system crash. In step 1108, the value of 10 is written to the parity disk at time T_D , thus returning the computed value of D1 to 1, which is correct. At time T_E , NV-RAM is cleared in step 1113. Thus, the present invention prevents data from being corrupted by a system fault when a disk is malfunctioning by using NV-RAM.

[0075] Figure 7B is a timing diagram of the present invention for preventing data corruption by storing data from a malfunctioning disk in NV-RAM for the case where parity is written to disk before data is updated. At time T_A , broken data disk 1 is illustrated having a value of 7 indicated within parentheses. This value is computed as described above with reference to Figure 7A. In step 1109, the stripe is read from the RAID array at time T_A . The NV-RAM is cleared which is indicated by a question mark under the heading for NV-RAM.

[0076] At time T_B , a value of 7 for the malfunctioning data disk 1 is written into NV-RAM according to step 1109. The value of 7 for data disk 1 that is written into NV-RAM is indicated by a rectangular box in Figure 7B. The system then computes a new value of parity for the stripe in step 1110 of Figure 11.

[0077] At time T_C , a new value of 10 (indicated by a box) for parity is written to the parity disk in step 1108 before data block 0 is updated. Therefore, at time T_C , the value for data disk 1 is 5 and is indicated within parentheses accordingly. When a system failure occurs between times T_C and T_D , a new parity value is updated correctly for the parity disk using NV-RAM. At decision block 1101 after the system reboots, a check is made if a system failure occurred. The decision block accordingly returns true (Yes) in the present example, and continues at step 1104.

[0078] Parity is correctly updated at time T_D by recalculating its value as follows:

$$\text{Parity} = \text{NV-data for broken disk (7)} + \text{on-disk data for all non broken}$$

$$\text{disks} = 4 + 7 + 1 = 12. \quad (13)$$

[0079] Thus, as shown in Figures 7A-7B, when the system is about to write to a stripe, it saves the value for malfunctioning data disk 1 in non-volatile RAM. It then writes the new value for data disk 0 (parity) to disk. If a system crash occurs after the new value is written to disk 0 (the parity disk) at time T_C , the value for data disk 1 is corrupt. After the system failure, the new value of parity (data disk 0) is calculated using the value of 7 stored in NV-RAM instead of the computed value of 5 for data disk 1. The value of parity (data disk 0) is then written to disk. Once this is completed, NV-RAM is erased.

Simultaneous System and Disk Failure Using NV-RAM

[0080] The present invention solves the problem of parity and data corruption when simultaneous system and disk failures occur by saving blocks of stripes in NV-RAM. Using NV-RAM allows the system to be recovered to a consistent state when a system crash occurs while updating multiple blocks (in the following example, data blocks 0 and 1) in the system. Changing these data blocks further requires that the parity of the stripe be updated. The present invention always saves into NV-RAM any block that is read from disk (e.g., before updating data block 0, read it into NV-RAM) for this purpose. Thus, stripe information can be recomputed from the data stored in NV-RAM. The present invention provides two solutions for this using parity by subtraction and parity by recalculation.

[0081] In parity by subtraction, data including parity and data blocks is read from disk before it is updated. Figure 10A is a timing diagram of the present invention for preventing parity and data corruption by storing blocks of a stripe in NV-RAM. The drawing is discussed with reference to Figures 11A-C. In Figure 10A, data is written to disk before parity is updated. At time T_A , the parity block and data block 0 are read from the RAID array. The NV-RAM is erased. This is indicated in Figure 10A by a question mark under the heading for NV-RAM.

[0082] At time T_B , the parity block and data block 0 are written into NV-RAM as they are read from disk. The parity block and data block 0 that are written into NV-RAM are indicated by a rectangular box in Figure 10A. The system then computes a new value of parity for a value of 2 for data block 0.

[0083] At time T_C , the new value of 2 (indicated by a box) for data disk 0 is written to the disk before parity for the stripe is updated. When a system failure occurs between times T_C and T_D , a disk in the RAID array malfunctions, and thus the present invention provides solutions for the three cases of a broken disk: the parity disk; data disk 0; and, data disk 2 (or 3). At decision block 1101, a check is made if a system failure occurred. The decision block accordingly returns true (Yes) in the present example, and continues at step 1104. The three cases of a broken disk due to system failure where parity is calculated by subtraction are shown in Figures 10A-10C, respectively.

[0084] At time T_D in Figure 10A, the parity disk malfunctions due to the system failure between times T_C and T_D . In this case, there is nothing to be done. No data is lost, and no information can be written to the parity disk.

[0085] Referring to Figure 10B, at time T_D , data disk 0 malfunctions due to the system failure between times T_C and T_D . The general equation for recalculating parity in this case is:

$$\begin{aligned} \text{parity} &= \text{"NV-value for broken disk"} \\ &+ \text{"on-disk values for all non-broken disks"} \end{aligned}$$

[0086] In the present example that becomes:

$$\text{parity} = \text{NV(Data 0)} + \text{Data 1} + \text{Data 2} = 4 + 7 + 1 = 12$$

[0087] In effect, the parity is being updated so as to restore the broken disk to the value stored for it in the NV-RAM. In this particular example, the new value for parity happens to match the old value. If other data blocks besides data 0 were also being updated, and if one of them reached disk before the system failure, then the new parity value would not match the old.

[0088] Referring to Figure 10C, at time T_D , data disk 1 malfunctions due to the system failure between times T_C and T_D . This case is handled by reading the parity and data disk 0 values from NV-RAM in step 1114 and writing them to disk (indicated by a box). Thus, the change to data disk 0 is overwritten, but the stripe is returned to a consistent state. Data disk 1 is indicated as being broken at times T_D and T_E by enclosing its value in parentheses. Thus, the value of broken data disk 1 is correctly computed to be 7 (indicated by a box) at time T_D as follows:

$$\text{Data 1} = \text{NV}(\text{Parity}) - \text{NV}(\text{Data 0}) - \text{Data 2} = 12 - 4 - 1 = 7, \quad (14)$$

where NV(Parity) and NV(Data 0) are the values for parity and data block 0 stored in NV-RAM. At time T_E , NV-RAM is cleared. Thus, in Figure 10C, the stripe is maintained in a consistent state by the present invention although data disk 1 (a broken disk 2 is handled similarly) malfunctions after the system failure between times T_C and T_D .

[0089] This case can also be addressed by first calculating the old contents of the broken disk as follows:

$$\text{D1-calc} = \text{NV-parity}$$

- "NV values for disks being updated".

- "on-disk values of data disks not being updated".

[0090] A new parity value is calculated based on:

$$\text{parity} = \text{"D1-calc from step above"}$$

+ on-disk values for all non-busted data disks".

Simultaneous System and Disk Failure with Parity by Recalculation

[0091] In parity by recalculation, the data blocks that are not being updated are first read from disk, and then parity is recalculated based on these values combined with the new data about to be written. This is typically used in cases where multiple data blocks are being updated at once, because it is more efficient than parity by subtraction in those cases. For simplicity, in the present example, only one block is updated. The techniques shown apply for updates of any number of blocks.

[0092] Figures 12A-12C are timing diagrams for parity by recalculation. For times T_A , T_B and T_C they are all identical.

[0093] At time T_A in step 1109, blocks D1 and D2 are read from disk. In step 1110, the system computes the new parity based on the new data for disk 0 along with the data just read from disks 1 and 2.

[0094] At time T_B in step 1111, blocks D1 and D2 are written into NV-RAM, along with an indication of the stripe to which they belong.

[0095] At time T_C , during step 1112, the new value "2" is written to disk 0. In the normal case, the parity block would also have been written during step 1112, and there would be no corruption.

[0096] In the present example, there is a system failure in combination with a disk failure. When the system reboots after a system failure, execution begins at step 1101. Because there is a failure, the decision block returns true (Yes) and continues at step 1102 and performs the necessary steps to recover the RAID sub-system based on the contents of NV-RAM.

[0097] Figure 12A shows the case where the parity disk fails. In this case, nothing needs to be done. There is no possibility of data loss, because no data disks have failed.

[0098] Figure 12B shows the case where the disk being updated fails. Note that at time T_D , the calculated value for disk 0 is incorrect. In general, if multiple blocks are being updated, there is not enough information in NV-RAM to reconstruct the lost data block. This is because with parity by recalculation, it is the data that is not being updated that is loaded into NV-RAM. The data on the failed disk is not saved anywhere.

[0099] In this case, the present invention computes a new parity value that sets the contents of the failed disk to zero. The general equation for this is:

$$\text{parity} = \text{sum of non-broken disks}$$

And in this example that is:

$$\text{parity} = \text{D1} + \text{D2} = 7 + 1 = 8$$

[0100] At time T_E , the new parity value is written, and at time T_F , the NV-RAM values for D1 and D2 are cleared.

[0101] With a prior-art file system that writes new data in the same location as old data, zeroing out a data block would be unacceptable. But with WAFL, which always writes new data to unused locations on disk, zeroing a block that was being written has no harmful effect, because the contents of the block were not part of the file system.

[0102] Figure 12C shows the case where the disk not being updated fails. Note that at time T_D , the calculated value for disk 1 is incorrect. The equation to recalculate parity is:

parity = "NV-RAM value for failed disk"

+ "on-disk values for non-failed disks:

In the present example, that is:

$$\text{parity} = \text{NV}(D1) + D0 + D2 = 7 + 2 + 1 = 10$$

[0103] At time T_E , the new parity value is written, and at time T_F , the NV-RAM values for D1 and D2 are cleared.

[0104] In this manner, a method and apparatus are disclosed for providing error correction for an array of disks using non-volatile random access memory (NV-RAM).

Claims

1. A method for providing parity correction for a RAID array in a computer system after a system failure, the method including steps of:

maintaining (1111, 1113) information about stripes having possibly inconsistent parity during routine operation of the computer system;

after said system failure, identifying (1104) stripes in response to said information; and

for each said stripe so identified, correcting (1106-1108) said possibly inconsistent parity.

2. A method according to claim 1, wherein said step of identifying includes the step of rebooting.

3. A method according to claim 1 or 2, wherein said information includes stripe number information.

4. A method according to claim 1, 2 or 3, wherein said step of correcting includes steps of:

reading (1106) data blocks of said each said stripe necessary to recompute a correct parity for said each stripe;

recomputing (1107) said correct parity;

writing (1108) said correct parity to said stripe.

5. A method according to any one of the preceding claims, wherein said step of maintaining includes steps of:

when updating data, a portion of said data destined for a first one or more blocks of a stripe having a stripe number and an existing parity, reading (1109) a second one or more blocks of said stripe required to compute a new parity of said stripe;

computing (1110) said new parity for said stripe;

adding (1111) said stripe number to a list of dirty stripes;

writing (1112) said portion of said data to said stripe;

writing said new parity to said stripe;

removing (1113) said stripe number from said list of dirty stripes.

6. A method according to claim 5, wherein said second one or more blocks includes said existing parity block.

7. A method according to claim 5, wherein said second one or more blocks excludes said existing parity block.

8. A method according to any one of the preceding claims, wherein said information is stored on non-volatile RAM (816).

9. A method according to any one of the preceding claims wherein said information includes a list of dirty stripes.
10. A method according to any one of the preceding claims, wherein said information includes a bitmap and including the further step of, when said bitmap is sufficiently full and after all pending writes and successfully completed:
clearing said bitmap.
11. A method according to claim 10, wherein each bit in said bitmap is associated with a group of one or more stripes.
12. A method according to claim 11, wherein each said group is 32 adjacent stripes.
13. A method according to claim 11, wherein said step of maintaining includes steps of:
 - when updating data, a portion of said data destined for a first one or more blocks of a stripe having an associated group and an existing parity, reading a second one or more blocks of said stripe used to compute a new parity of said stripe;
 - computing said new parity for said stripe;
 - selecting a bit in said bitmap responsive to said group;
 - setting said bit;
 - writing said portion of said data to said stripe;
 - writing said new parity to said stripe.
14. A method according to claim 10, wherein said RAID array includes one or more write-caching disks, each said one or more write-caching disks having an internal cache and including the further step of, prior to said step of clearing said bitmap:
 - flushing each said internal cache.
15. An apparatus for providing parity correction for a RAID array in a computer system after a system failure, said apparatus includes:
 - a maintenance mechanism (1111-1113) configured to maintain information about stripes having possibly inconsistent parity during routine operation of the computer system;
 - an identification mechanism (1104) configured to identify stripes in response to said information, said identification apparatus activated in response to check on reboot after a system failure;
 - a correction mechanism (1106-1108) configured to correct said possibility inconsistent parity for each of said stripes identified by said identification mechanism.
16. An apparatus according to claim 15 wherein said information includes stripe number information.
17. The apparatus according to claim 15 or 16 wherein correction mechanism includes:
 - a read mechanism configured mechanism configured to recompute said correct parity responsive to said reading;
 - a write mechanism configured to write said correct parity to said stripe responsive to said recomputing.
18. An apparatus according to claim 15, 16 and 17 wherein said maintenance mechanism further includes:
 - a read mechanism (1109) configured to read a second one or more blocks of said stripe required to compute a new parity of said stripe, said mechanism activated when new data is to be written to a first one or more blocks of a stripe having a stripe number and existing parity;
 - a compute mechanism (1110) configured to compute said new parity for said stripe responsive to said reading;
 - an add mechanism (1111) configured to add said stripe number to a list of dirty stripes responsive to said compute mechanism;
 - a write mechanism (1112) configured to write said data to said stripe;
 - a second write mechanism configured to write said new parity to said stripe;
 - a remove mechanism (1113) configured to remove said stripe number from said list of dirty stripes.
19. An apparatus according to claim 18 wherein said second one or more blocks includes said existing parity block.

20. An apparatus according to claim 18 wherein said second one or more blocks excludes said existing parity block.
21. An apparatus according to any one of claims 15 to 20 wherein said information is stored on non-volatile RAM (816).
- 5 22. An apparatus according to any one of claims 15 to 21 wherein said information includes a list of dirty stripes.
23. An apparatus according to any one of claims 15 to 22 wherein said information includes a bitmap and further including:
 - 10 a clear mechanism configured to clear said bitmap when said bitmap is sufficiently full and after all pending writes are successfully completed:
24. An apparatus according to claim 23 wherein each bit in said bitmap is associated with a group of one or more stripes.
- 15 25. An apparatus according to claim 23 or 24 wherein said RAID array includes one or more write-caching disks, each said one or more write-caching disks having an internal cache and further including:
 - a flush mechanism configured to flush each said internal cache.
26. An apparatus according to any one of claims 15 to 25 wherein said maintenance mechanism includes:
 - 20 a read mechanism (1106) configured to read a first one or more blocks of said stripe used to compute a new parity of stripe, a second one or more blocks of said stripe to be updated with new data, said stripe having an associated group and an associated parity block;
 - a second compute mechanism configured to compute said new parity for said stripe, responsive to said read mechanism;
 - 25 a selector configured to select a bit in said bitmap responsive to said group;
 - a set mechanism configured to set said bit;
 - a first write mechanism configured to write new data to said stripe;
 - a second write mechanism configured to write said new parity to said stripe.
- 30 27. A computer program comprising computer program code means adapted to perform the steps of any one of claims 1 to 14 when said program is put on a computer.
28. A computer program according to claim 27 embodied on a computer readable recording medium.
- 35 29. A method for maintaining a consistent file system, wherein said file system is stored in stripes of blocks in a RAID array, that allows for recovery of the file system to a consistent state after a disk failure and a system failure occur prior to successful completion of writing new data to one or more blocks of a stripe of said RAID array, said disk failure rendering a block of said stripe in said RAID array unreadable, the method comprising the steps of:
 - 40 prior to writing new data to one or more blocks of a stripe in said RAID array, reading said one or more blocks of said stripe to obtain prior data stored in said one or more blocks of said stripe, said prior data stored in said stripe storage means comprising a prior parity value read from a parity block of said stripe and prior data read from one or more data blocks of said stripe;
 - storing said prior data stored in said one or more blocks of said stripe in stripe storage means;
 - 45 initiating the writing of said new data to said one or more blocks of said stripe; and
 - if said writing of said new data to said one or more blocks of said stripe in said RAID array is successfully completed, clearing said prior data from said one or more blocks of said stripe storage means; otherwise, performing the steps of;
 - after a failure has occurred, checking said stripe storage means to determine whether said stripe storage means contains any prior data for any stripe;
 - 50 when said stripe storage means contains prior data for any stripe, performing the steps of:
 - determining whether said block of said stripe in said RAID array rendered unreadable by said disk failure is one of said data blocks of said stripe in said RAID array for which prior data has been stored in said stripe storage means;
 - 55 if said block of said stripe in said RAID array rendered unreadable by said disk failure is one of said data blocks of said stripe in said RAID array for which prior data has been stored in said stripe storage means, performing the steps of:
 - reading data from each remaining readable data block of said stripe in said RAID array other than said data

block rendered unreadable by said disk failure;
 recomputing a parity value of said stripe based on said prior data stored in said stripe storage means for said
 data block rendered unreadable by said disk failure and on said data read from said each remaining readable
 data block of said stripe; and
 5 writing said recomputed parity value for said stripe to said parity block of said stripe in said RAID array.

30. A method for maintaining a consistent file system, wherein said file system is stored in stripe blocks in a RAID
 array, that allows for recovery of the file system to a consistent state after disk failure and system failure occur
 prior to completion of writing new data to one or more blocks of a stripe of said RAID array, said disk failure rendering
 10 a block of said stripe in said RAID array unreadable, comprising the steps of:

prior to writing new data to one or more blocks of stripe in said RAID array, reading said one or more blocks
 of said stripe to obtain prior data stored in said one or more blocks of said stripe, said prior data stored in said
 stripe storage means comprising a prior parity value read from a parity block of said stripe and prior data read
 15 from one or more data blocks of said stripe;
 storing said prior data stored in said one or more blocks of said stripe in stripe storage means;
 initiating the writing of said new data to said one or more blocks of said stripe; and
 if said writing of said new data to said one or more blocks of said stripe in said RAID array is successfully
 completed, clearing said prior data from said one or more blocks of said stripe from said stripe storage means;
 20 otherwise, performing the steps of:
 reading data from each remaining readable data block of said stripe in said RAID array other than said one
 or more data blocks of said stripe for which prior data has been stored in said stripe storage means;
 computing data for said data block rendered unreadable by said disk failure based on said data read from said
 each remaining data block of said stripe in said RAID array other than said one or more data blocks of said
 25 RAID array for which prior data has been stored in said stripe storage means and on said prior data or said
 one or more data blocks of said stripe and said prior parity value for said parity block of said stripe stored in
 said stripe storage means;
 reading data from said data blocks of said stripe in said RAID array for which prior data has been stored in
 said stripe storage means;
 30 recomputing a new parity value based on said computed data for said data block rendered unreadable by said
 disk failure, said data read from said data blocks of said stripe in said RAID array for which prior data has
 been stored in said stripe storage means; and said data read from said each remaining readable data block
 of said stripe in said RAID array other than said one or more data blocks for which prior data has been stored
 in said stripe storage means; and
 35 writing said new parity value for said stripe to said parity block of said stripe in said RAID array.

31. A method for maintaining a consistent file system, wherein said file system is stored in stripes of blocks in a RAID
 array, that allows for recovery of the file system to a consistent state after disk failure and system failure occur
 prior to completion of writing new data to one or more blocks of a stripe of said RAID array, said disk failure rendering
 40 a block of said stripe in said RAID array unreadable, comprising the steps of:

prior to writing new data to one or more blocks of a stripe in said RAID array, reading said one or more blocks
 of said stripe to obtain prior data stored in said one or more blocks of said stripe, said prior data stored in said
 stripe storage means comprising a prior parity value read from a parity block of said stripe and prior data read
 45 from one or more data blocks of said stripe;
 storing said prior data stored in said one or more blocks of said stripe in stripe storage means;
 initiating the writing of said new data to said one or more blocks of said stripe; and
 if said writing of said new data to said one or more blocks of said stripe in said RAID array is successfully
 completed, clearing said prior data from said one or more blocks of said stripe from said stripe storage means;
 50 otherwise, performing the steps of:
 after a failure has occurred, checking said stripe storage means to determine whether said stripe storage
 means contains any prior data for any stripe;
 when said stripe storage means contains prior data for any stripe, performing the steps of:
 determining whether said block of said stripe in said RAID array rendered unreadable by said disk failure is
 55 one of said data blocks of said stripe in said RAID array other than said one or more data blocks for which
 prior data has been stored in said stripe storage means;
 if said block of said stripe in said RAID array rendered unreadable by said disk failure is one of said data blocks
 of said stripe in said RAID array other than said one or more data blocks for which prior data has been stored

in said stripe storage means, writing said prior data for said one or more data blocks of said stripe and said prior parity value stored in said stripe storage means, writing said prior data for said one or more data blocks of said stripe and said prior parity value stored in said stripe storage means to said stripe in said RAID array.

- 5 **32.** A method for maintaining a consistent file system, wherein said file system is stored in stripes of blocks of a RAID array, that allows for recovery of the file system to a consistent state after a disk failure and a system failure occur prior to successful completion of writing new data to one or more blocks of a stripe of said RAID array, said disk failure rendering a block of said stripe in said RAID array unreadable, comprising the steps of:

10 prior to writing new data to one or more blocks of a stripe in said RAID array, reading any other blocks of said stripe in said RAID array other than said one or more blocks to which new data is to be written to obtain prior data stored in said other stripe;
 storing said prior data stored in said other blocks of said stripe in stripe storage means;
 initiating the writing of said new data to said one or more blocks of said stripe in said RAID array; and
 15 if said writing of said new data to said one or more blocks of said stripe is successfully completed, clearing said prior data from said blocks of said stripe from said stripe storage means; otherwise performing the steps of:
 after a failure has occurred, checking said stripe storage means to determine whether said stripe storage means contains any prior data for any stripe;
 when said stripe storage means contains prior data for any stripe, performing the steps of:
 20 determining whether said block of said stripe in said RAID array rendered unreadable by said disk failure is one of said data blocks of said stripe in said RAID array for which prior data has been stored in said stripe storage means;
 if said block of said stripe in said RAID array rendered unreadable by said disk failure is one of said data blocks of said stripe in said RAID array for which prior data has been stored in said stripe storage means, performing
 25 the steps of:
 reading data from each remaining readable data block of said stripe in said RAID array other than said data block rendered unreadable by said disk failure;
 recomputing a new parity value for said stripe based upon said data read from said each remaining readable data block of said stripe in said RAID array and said prior data stored in said stripe storage means for said
 30 data block rendered unreadable by said disk failure; and
 writing said new parity value for said stripe to a parity block of said stripe in said RAID array.

33. A method for operating a file system that provides protection against corruption of the file system upon writing new data to the file system after a disk failure has rendered unreadable a data block in a stripe of RAID array wherein
 35 file system is stored, and that allows recovery of the file system to a consistent state when a system failure occurs prior to all of said new data being successfully written and comprising the steps of:

 prior to writing new data to one or more blocks of said stripe, reading all readable blocks of said stripe in said RAID array;
 40 computing data for said unreadable block based upon said readable blocks;
 storing said computed data in stripe storage means;
 initiating the writing of said new data to said one or more blocks of said stripe in said RAID array; and
 if said writing of said new data to said one or more blocks of said stripe is successfully completed, clearing
 45 said computed data from said stripe storage means; otherwise, performing the steps of:
 after said system failure has occurred, checking said stripe storage means to determine whether said stripe storage means contains any computed data for any unreadable data block of any stripe;
 when said stripe storage means containing computed data for an unreadable data block of a stripe, performing
 the steps of:
 reading data from each readable data block of said stripe failure in said RAID array other than said data block
 50 rendered unreadable by said disk;
 recomputing a parity value for said stripe based on said computed data for said unreadable data blocks of said stripe stored in said stripe storage means and said data read from said remaining readable data blocks
 in said stripe in said RAID array; and
 writing said recomputed parity value for said stripe to a parity block of said stripe in said RAID array.
 55




FIG. 1
PRIOR ART

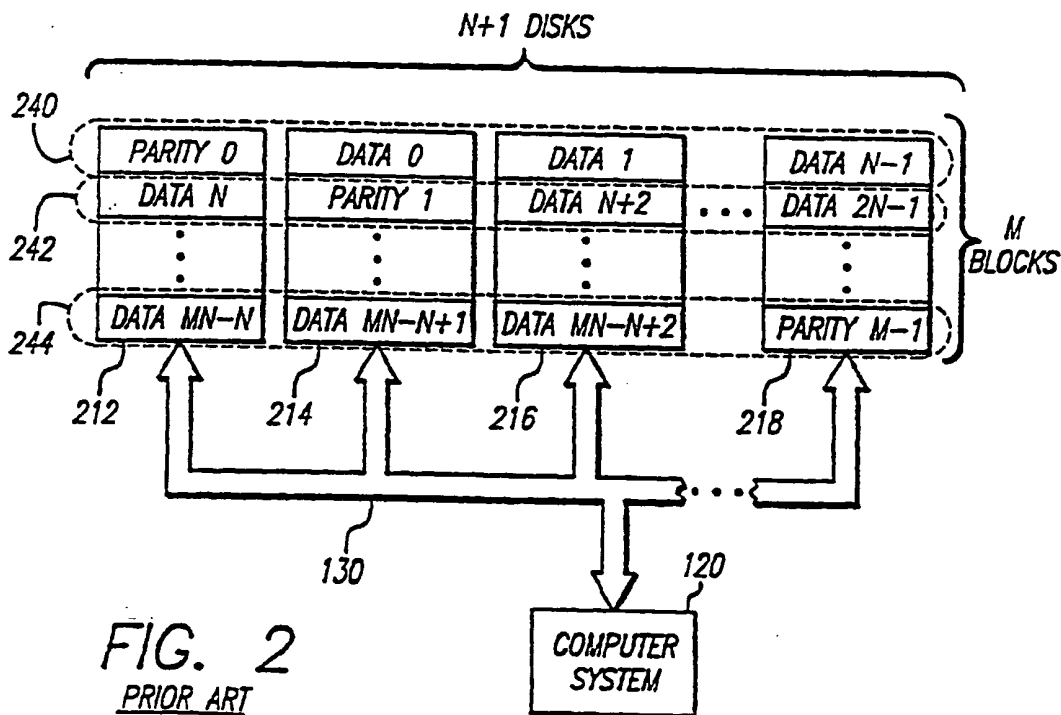


FIG. 2
PRIOR ART




FIG. 3A
PRIOR ART

0	<u>PARITY</u>	<u>DATA 0</u>	<u>DATA 1</u>	<u>DATA 2</u>
t_A	12	4	7	1
t_B	12	4	?	1
t_C	12	4	(7)	1

TIME

FIG. 3B
PRIOR ART




FIG. 3C
PRIOR ART

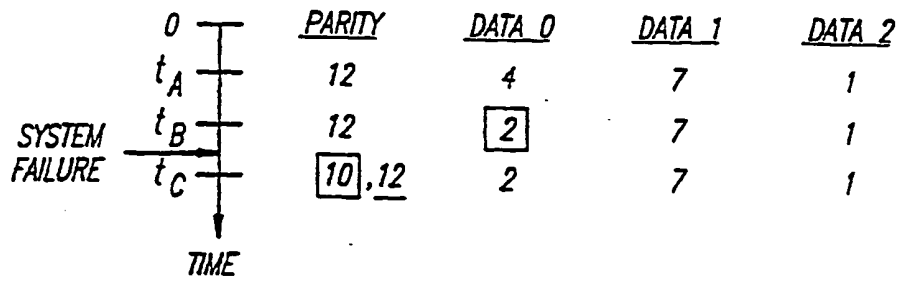


FIG. 4A

PRIOR ART

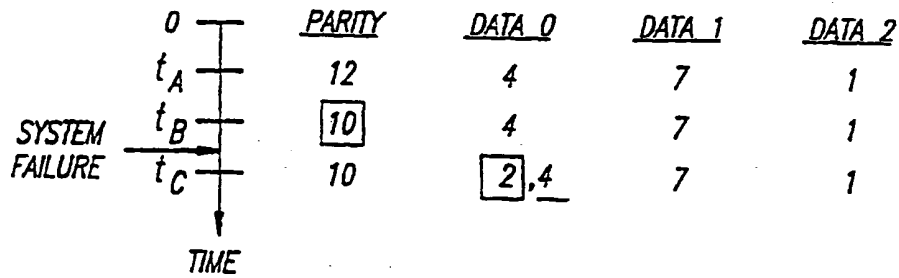


FIG. 4B

PRIOR ART

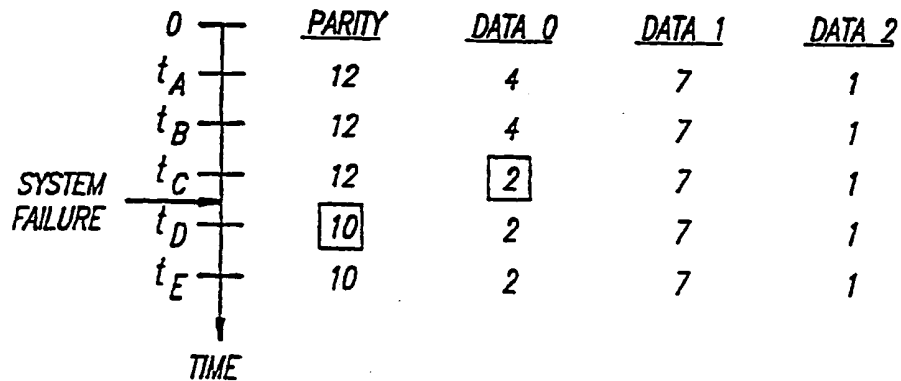


FIG. 5

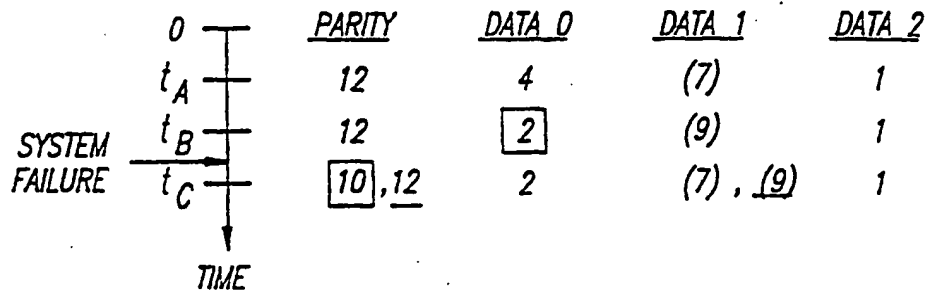


FIG. 6
PRIOR ART

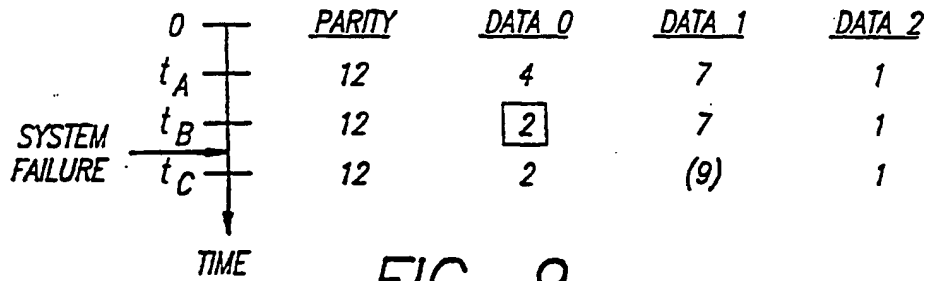


FIG. 9 PRIOR ART

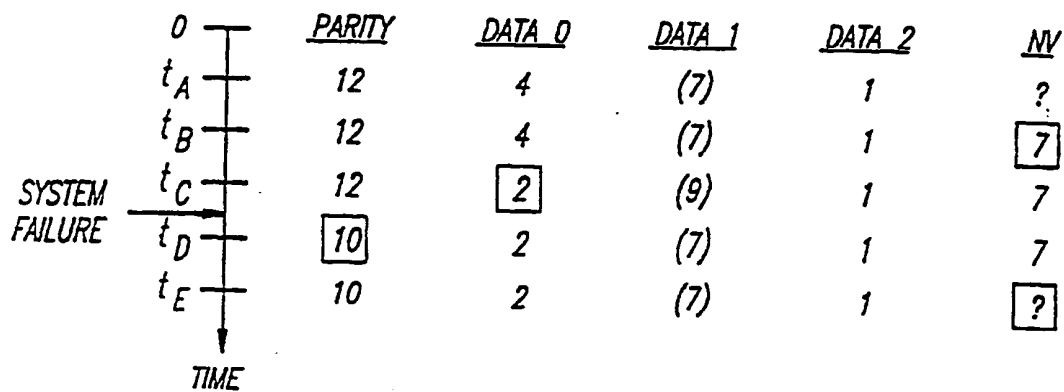


FIG. 7A

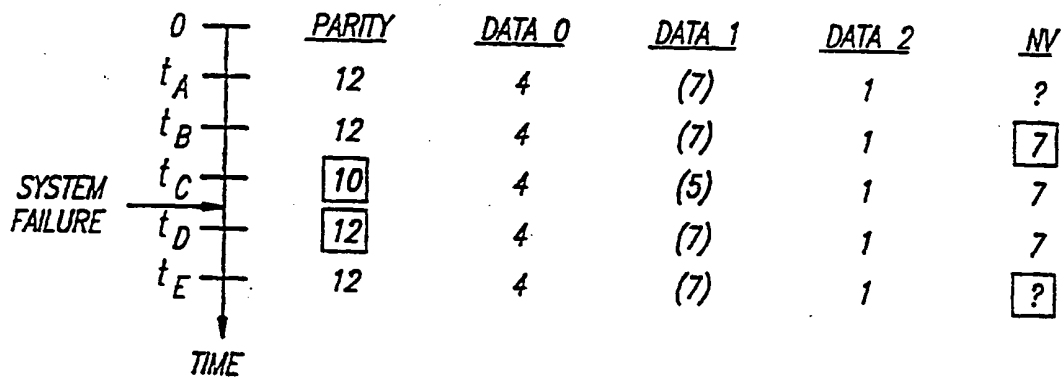


FIG. 7B




FIG. 8

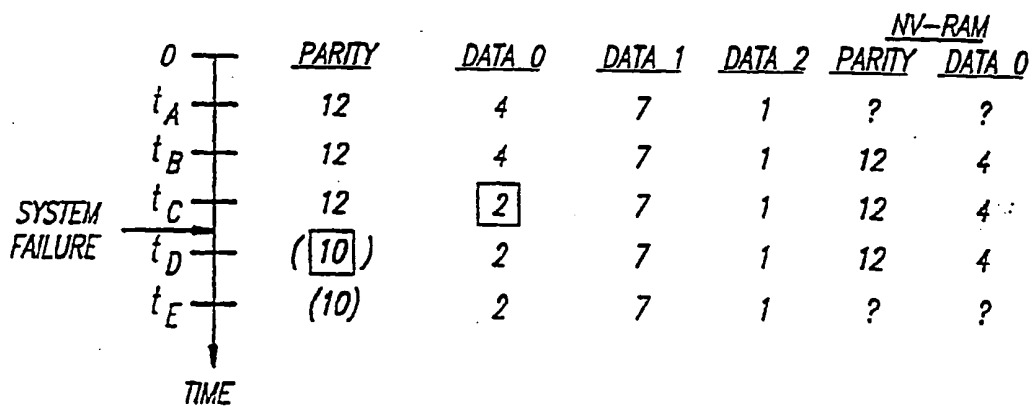


FIG. 10A

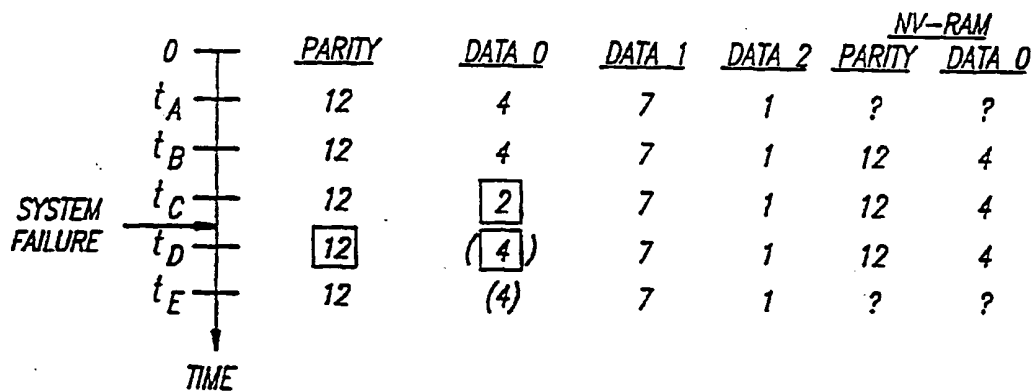


FIG. 10B

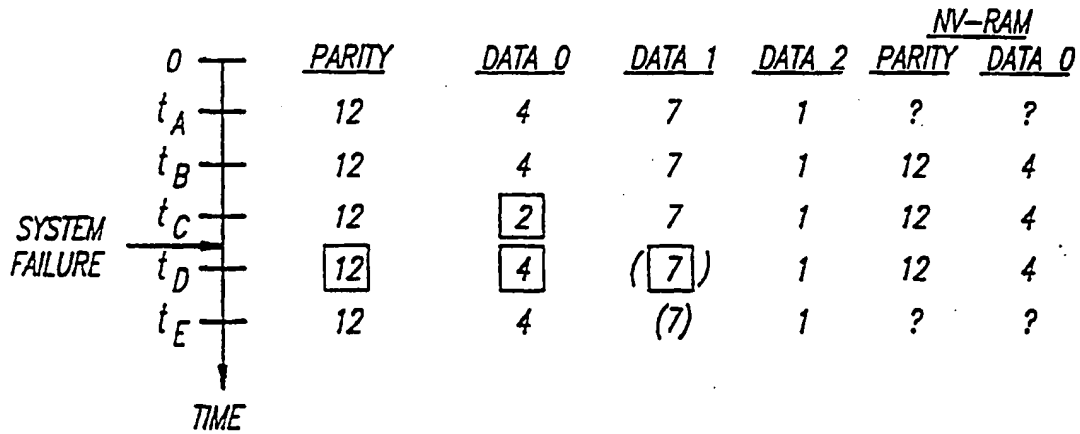


FIG. 10C






FIG. 11A





					NV-RAM	
	PAR	DO	D1	D2	D1	D2
T_D	12	4	7	1	?	?
T_B	12	4	7	1	171	111
T_C	12	121	7	1	7	1

T_D	(10)	2	7	1	7	1
T_E	(10)	2	7	1	1?1	1?1

FIG. 12A

					NV-RAM	
	PAR	DO	D1	D2	D1	D2
T_A	12	4	7	1	?	?
T_B	12	4	7	1	171	111
T_C	12	121	7	1	7	1

T_D	12	(4)	7	1	7	1
T_E	181	(0)	7	1	7	1
T_F	8	(0)	7	1	1?1	1?1

5

FIG. 12B

					<i>NV-RAM</i>	
	<i>PAR</i>	<i>D0</i>	<i>D1</i>	<i>D2</i>	<i>D1</i>	<i>D2</i>
T_A	12	4	7	1	?	?
T_B	12	4	7	1	171	111
T_C	12	121	7	1	7	1

T_D	12	2	(9)	1	7	1
T_E	1101	2	(7)	1	7	1
T_F	10	2	(7)	1	1?1	1?1

10

FIG. 12C

